arXiv:1510.06051v1 [math.CO] 20 Oct 2015

THE COMPLEXITY OF PATTERN MATCHING FOR 321-AVOIDING AND SKEW-MERGED PERMUTATIONS

Michael H. Albert
Department of Computer Science
University of Otago
Dunedin, New Zealand

Marie-Louise Lackner*
Institute of Discrete Mathematics and Geometry
Vienna University of Technology
Vienna, Austria

Martin Lackner†
Institute of Information Systems
Vienna University of Technology
Vienna, Austria

Vincent Vatter‡
Department of Mathematics
University of Florida
Gainesville, Florida USA

The PERMUTATION PATTERN MATCHING problem, asking whether a pattern permutation π is contained in a permutation τ, is known to be NP-complete. In this paper we present two polynomial time algorithms for special cases. The first algorithm is applicable if both π and τ are 321-avoiding; the second is applicable if π and τ are skew-merged. Both algorithms have a runtime of $O(kn)$, where k is the length of π and n the length of τ.

1. INTRODUCTION

In this paper, a permutation is a bijective function from $[n]$ to itself, where n is a positive integer and $[n] = \{1, 2, \ldots, n\}$. Therefore, a permutation, $\pi : [n] \to [n]$ *is* the set of ordered pairs $(i, \pi(i))$. The *size* of π is just the cardinality of this set, and we denote the elements, also called points, of a permutation by variables such as x and y. We adopt the usual conventions with respect to order of such points, i.e., $(i, \pi(i))$ lies to the left of $(j, \pi(j))$ if $i < j$, and above $(k, \pi(k))$ if $\pi(i) > \pi(k)$ (with corresponding definitions for 'to the right of' and 'below'). Given an element x in a permutation π, we define (wherever possible):

*The author's research was supported by the Austrian Science Foundation FWF, grant P25337-N23.
†The author's research was supported by the Austrian Science Foundation FWF, grants P25518-N23 and Y698.
‡The author's research was partially supported by the National Science Foundation under Grant Number DMS-1301692.
Date: October 22, 2015
AMS 2010 Subject Classification. 05A05, 68Q25

$x^{\blacktriangleleft}$ the element immediately to its left,
$x^{\blacktriangleright}$ the element immediately to its right,
$x^{\blacktriangle}$ the element immediately above it, and
$x^{\blacktriangledown}$ the element immediately below it.

We use the symbol $\perp$ to represent 'undefined' and so for instance, $x^{\blacktriangleleft} = \perp$ means that x is the leftmost element of π. Any operator applied to $\perp$ also yields $\perp$. We occasionally write specific permutations in the usual one line notation, e.g., 321 represents the permutation of [3] equal to $\{(1,3),(2,2),(3,1)\}$.

Let π and τ be two permutations. An injective function, f, from π into τ is an *embedding* if, for all elements x and y of π, the elements $f(x)$ and $f(y)$ of τ are in the same relative order as x and y (e.g., if x lies below and to the right of y, then $f(x)$ also lies below and to the right of $f(y)$). If there is an embedding from π into τ then we say that τ *contains* π. If not, then we say that τ *avoids* π. The following problem, analogous to problems related to detecting the occurrence of a 'pattern' in a 'text', is central with respect to these concepts:

PERMUTATION PATTERN MATCHING (PPM)	
Instance:	A text permutation τ of size n and a pattern π of size k.
Question:	Does τ contain π?

Bose, Buss, and Lubiw [5] showed in 1998 that the PPM problem is NP-complete[1].

The recent work of Guillemot and Marx [9] shows that the PPM problem can be solved in time $2^{O(k^2 \log k)}n$, i.e., linear-time in n when k is viewed as a constant. In particular, this implies that the PPM problem is *fixed-parameter tractable (fpt)* with respect to the length of the pattern k. Work prior to this breakthrough result achieved runtimes of $O(n^{1+2k/3} \cdot \log n)$ [2] and $O(n^{0.47k+o(k)})$ [1]. Using $\mathsf{run}(\tau)$ as parameter, i.e., the number of alternating runs of τ, Bruner and Lackner [6] present an fpt algorithm with runtime $O(1.79^{\mathsf{run}(\tau)} \cdot kn)$.

A *permutation class*, $\mathcal{C}$, is a set of permutations with the property that, if $\tau \in \mathcal{C}$ and τ contains π, then $\pi \in \mathcal{C}$. In other words, $\mathcal{C}$ is closed downwards with respect to the partial ordering of permutations given by the relation "is contained in". A permutation class is proper if it does not contain every permutation. A permutation class is frequently defined in terms of avoidance conditions, namely, for any set B of permutations, the set $\mathrm{Av}(B)$ consisting of those permutations which avoid (do not contain) every element of B is a permutation class (and it is always proper if B is non-empty). Every permutation class is equal to the class of permutations avoiding its complement, or even avoiding the minimal elements (with respect to the partial ordering mentioned above) of its complement.

This leads to two natural ways in which one might restrict the PPM problem. One is to impose additional structure on the pattern, most naturally, to insist that the pattern belongs to a particular (proper) permutation class. One example which has been studied is the class of *separable* permutations, which are those avoiding both 3142 and 2413. If the pattern is separable, PPM can be solved in polynomial time [2, 5, 11, 12]. The fastest algorithm for this case is by Ibarra [11] with a runtime of $O(kn^4)$. Formally, we define this class of problems as follows:

[1] Note that because both the pattern and text are regarded as input, the size of the input is $n + k$. Were we to regard the size of the pattern as fixed, then the trivial $\binom{n}{k}$ algorithm would be polynomial time.

$\mathcal{C}$-PATTERN PERMUTATION PATTERN MATCHING ($\mathcal{C}$-pattern PPM)	
Instance:	A text permutation τ of size n and a pattern π of size k, where π belongs to a fixed proper permutation class $\mathcal{C}$.
Question:	Does τ contain π?

A second and more restrictive specialisation of the PPM problem is to insist that both the pattern and text belong to a (proper) permutation class. This is the version of the problem that we study.

$\mathcal{C}$ PERMUTATION PATTERN MATCHING ($\mathcal{C}$-PPM)	
Instance:	A text permutation τ of size n and a pattern π of size k, both belonging to a fixed proper permutation class $\mathcal{C}$.
Question:	Does τ contain π?

Clearly, for a fixed $\mathcal{C}$, polynomial time algorithms for $\mathcal{C}$-PATTERN PPM apply to $\mathcal{C}$-PPM as well. Consequently, the separable case, i.e., Av(3142, 2413)-PPM, can be solved in $O(kn^4)$ time [11]. Most relevant to our work is a result by Guillemot and Vialette [10] that establishes an $O(k^2n^6)$-time algorithm for Av(321)-PPM. In Sections 2 and 3, we improve their approach to give the following.

Theorem 1. *Given* 321*-avoiding permutations τ of size n and π of size k, there is an $O(kn)$-time algorithm which determines whether τ contains π.*

Then in Section 4 we show how to adapt this approach to the class of skew-merged permutations, which are those permutations whose elements can be partitioned into an increasing subsequence and a decreasing subsequence. Skew-merged permutations can also be characterized as those permutations that avoid both 3412 and 2143 [13].

Theorem 2. *Given skew-merged permutations τ of size n and π of size k, there is an $O(kn)$-time algorithm which determines whether τ contains π.*

The following elementary observation will be used repeatedly.

Observation 3. *Let π and τ be permutations and $f : \pi \to \tau$. Then f is an embedding of π into τ if and only if for every element x of π:*

- *if $x^{\blacktriangleleft} \neq \perp$ then $f(x)$ lies strictly to the right of $f(x^{\blacktriangleleft})$ and*
- *if $x^{\blacktriangledown} \neq \perp$ then $f(x)$ lies strictly above $f(x^{\blacktriangledown})$.*

Proof. Suppose that x and y are points of π and that, without loss of generality, y lies to the left of x. Then y occurs in the sequence $x, x^{\blacktriangleleft}, x^{\blacktriangleleft\blacktriangleleft}, x^{\blacktriangleleft\blacktriangleleft\blacktriangleleft}, \ldots$. So, by inductive use of the first property, $f(x)$ lies to the right of $f(y)$. Similarly, inductive use of the second property establishes that the vertical relationship between $f(x)$ and $f(y)$ is the same as that between x and y, and the result follows. □

2. THE LATTICE OF RIGID EMBEDDINGS OF 321-AVOIDING PERMUTATIONS

It is a well-known fact that the elements of any 321-avoiding permutation, π, can be partitioned into two increasing subsequences. In any such partition, one of these subsequences will contain

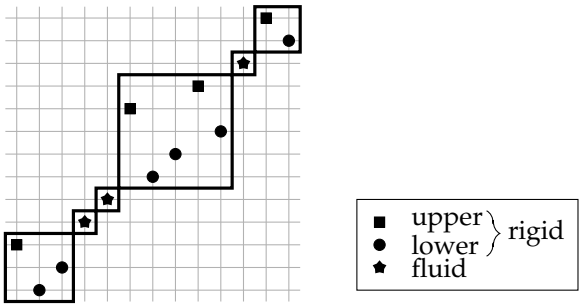

Figure 1: The decomposition of the 321-avoiding permutation $\pi = 3\,1\,2\,4\,5\,9\,6\,7\,10\,8\,11\,13\,12$ into rigid and fluid elements.

all those elements which participate as the '2' in a copy of 21—called the *upper elements* of π and denoted U_π— and the other will contain all those elements which participate as the '1' in a copy of 21—called the *lower elements* of π and denoted L_π. We hope that the meaning of these informal descriptions is clear, but just to be complete, an element x of π participates as a 2 in a copy of 21 if there is some embedding of $21 = \{(1,2),(2,1)\}$ into π such that x is the image of $(1,2)$.

Following Albert, Atkinson, Brignall, Ruškuc, Smith, and West [3], elements which are either upper or lower elements of π are referred to as *rigid* elements, and π is called a *rigid permutation* if all of its elements are rigid (i.e., if $\pi = U_\pi \cup L_\pi$). The remaining elements will be called *fluid* elements. An element is thus called fluid if it does not participate in a 21-pattern. For an example of a 321-avoiding permutation and its decomposition into rigid and fluid elements, see Figure 1.

The existence of fluid elements in a pattern will be the source of some difficulty in solving the Av(321)-PPM, and will be addressed in the next section. For the remainder of this section we fix a rigid pattern π of size k and a 321-avoiding text τ of size n. Since an embedding preserves relative locations of points and their images, the image of any rigid element must be rigid. Therefore, if there exists an embedding of π into τ, then it may only map upper (resp., lower) elements of π to upper (resp., lower) elements of τ; the fluid elements of τ will never be used.

In order to look for such embeddings we must widen our search space. A map $f : \pi \to \tau$ is called a *rigid mapping* if f maps upper (resp., lower) elements of π to upper (resp., lower) elements of τ. As noted above, because π is rigid, every embedding of π into τ is a rigid mapping, but the converse is far from true since, among other reasons, rigid mappings need not be injective.

Given two points, x and y, in U_π, we say $x \leqslant y$ if y lies above and to the right of x. This is a linear order on U_π, and we have similar linear orders (all denoted $\leqslant$) on L_π, U_τ and L_τ. This makes the set of all rigid mappings of π into τ into a partially ordered set using point-wise comparison; that is, given rigid mappings $f, g : \pi \to \tau$, we write $f \leqslant g$ if $f(x) \leqslant g(x)$ for all elements x of π. In fact, it is easy to see that this partially ordered set is a distributive lattice; given two rigid mappings $f, g : \pi \to \tau$ their *meet* and *join* can be defined, respectively, by

$$\begin{aligned}(f \wedge g)(x) &= \min\{f(x), g(x)\},\\ (f \vee g)(x) &= \max\{f(x), g(x)\}\end{aligned}$$

for all elements x of π.

It is notable that these observations also hold for embeddings. That is, the set of embeddings from π into τ is a sublattice of the set of rigid mappings:

Theorem 4 ([3, Theorem 2])**.** *Given a rigid pattern π and 321-avoiding text τ, the set of embeddings of π into τ forms a distributive lattice under the operations of meet and join defined above.*

It follows from Theorem 4 that if π is contained in τ then there is a *minimum embedding* of π into τ, which we denote $e_{\min}$.

Given an element x of some 321-avoiding permutation, σ, we define $x^{\blacktriangleleft}_U$ to be the rightmost element of U_σ that is to the left of x. Of course, we have corresponding notations such as $x^{\blacktriangledown}_L$, $x^{\blacktriangleright}_U$ and so on. In all cases, if no such element exists we get $\perp$ as usual. We also define the *type* of x, $T(x)$ to be U if $x \in U_\sigma$ and L if $x \in L_\sigma$. The following result forms the core of our algorithm for determining whether there is an embedding of π into τ, at least for the case where π is rigid. It will allow us to turn an arbitrary rigid mapping into an embedding, if possible.

Proposition 5. *Suppose that $e : \pi \to \tau$ is an embedding, $f : \pi \to \tau$ is a rigid mapping, and, for all $x \in \pi$, $f(x) \leqslant e(x)$. Then, for all $x \in \pi$:*

$$\max\{f(x^{\blacktriangleleft})^{\blacktriangleright}_{T(x)}, f(x^{\blacktriangledown})^{\blacktriangle}_{T(x)}\} \leqslant e(x).$$

Proof. We first establish that, $f(x^{\blacktriangleleft})^{\blacktriangleright}_{T(x)} \leqslant e(x)$. Since x lies to the right of $x^{\blacktriangleleft}$ (and e is an embedding), $e(x)$ lies to the right of $e(x^{\blacktriangleleft})$ and is of the same type as x, so, it does not lie to the left of $e(x^{\blacktriangleleft})^{\blacktriangleright}_{T(x)}$. That is, $e(x^{\blacktriangleleft})^{\blacktriangleright}_{T(x)} \leqslant e(x)$. But $f(x^{\blacktriangleleft}) \leqslant e(x^{\blacktriangleleft})$ and so $f(x^{\blacktriangleleft})^{\blacktriangleright}_{T(x)} \leqslant e(x^{\blacktriangleleft})^{\blacktriangleright}_{T(x)} \leqslant e(x)$. The arguments for the other case is exactly similar. □

Applying the proposition above in the case where $f = e$, we see that for any embedding, e, from a rigid π to τ, and any $x \in \pi$:

$$\max\{e(x^{\blacktriangleleft})^{\blacktriangleright}_{T(x)}, e(x^{\blacktriangledown})^{\blacktriangle}_{T(x)}\} \leqslant e(x).$$

Now suppose that f is any rigid mapping from π to τ. We say that x *is a problem* if it violates the above condition, i.e., if

$$f(x) < \max\{f(x^{\blacktriangleleft})^{\blacktriangleright}_{T(x)}, f(x^{\blacktriangledown})^{\blacktriangle}_{T(x)}\}.$$

Intuitively , x is a problem if $f(x)$ is too low compared with $f(x^{\blacktriangledown})$ or too far left compared with $f(x^{\blacktriangleleft})$. We let $P(f)$ be the set of problems for f. The discussion above together with Observation 3 together imply that f is an embedding if and only if $P(f)$ is empty.

We now describe an algorithm, displayed as Algorithm 1, that, given as input a rigid permutation π and a 321-avoiding permutation τ returns $e_{\min}$ the minimum embedding of π into τ when it exists, and fails otherwise. The algorithm constructs and updates a rigid mapping f, ensuring that $f \leqslant e_{\min}$ at all times (in the first case)[2]. Let f_0 be the map that sends all the elements of U_π to the least element of U_τ and all elements of L_π to the least element of L_τ.

The correctness of this algorithm is easy to establish. Within the while loop, if f is everywhere defined, $P(f)$ is non-empty, and x is chosen for the update step, then the updated version of f

[2]The mathematical purist may choose to believe that the algorithm constructs a sequence of rigid mappings $f_0, f_1, \ldots$, and terminates when it has constructed an embedding, or when it is impossible to define all the values of f_{k+1} from those of f_k

Algorithm 1 Find a minimum embedding of π into τ, or demonstrate that no embeddings exist.

Initialise: $f \leftarrow f_0$.
while f is defined everywhere, and $P(f)$ is non-empty **do**
  Choose $x \in P(f)$.
  Update: $f(x) \leftarrow \max\{f(x^{\blacktriangleleft})^{\blacktriangleright}_{T(x)}, f(x^{\blacktriangledown})^{\blacktriangle}_{T(x)}\}$
  Recompute: $P(f)$.
end while
Return: f, which, if everywhere defined, equals $e_{\min}$.

is strictly greater than the original at x, and has the same value elsewhere. Since the set of rigid maps is finite, the loop can be executed a bounded number of times, and the algorithm halts. In the case where $e_{\min}$ exists, we certainly have $f_0 \leqslant e_{\min}$. So, by Proposition 5, it is always the case that $f \leqslant e_{\min}$. Therefore, when the loop terminates, the algorithm returns an embedding that is less than or equal to $e_{\min}$, and hence must equal $e_{\min}$. Should $e_{\min}$ not exist, then termination can only occur because f is not everywhere defined, and so the algorithm fails as required in this case.

We can further combine the correctness analysis with a run-time analysis to obtain:

Proposition 6. *Given a rigid* 321*-avoiding permutation π of size k and a* 321*-avoiding permutation τ of size n there is an algorithm which determines an embedding of π into τ if one exists, and fails otherwise, whose run-time is $O(kn)$.*

Proof. The algorithm in question is Algorithm 1, and it remains to show that we can achieve the bound claimed for the run-time. As noted, each execution of the loop increases the value of $f(x)$ for at least one x (in the linear ordering, $\leqslant$, of either U_τ or L_τ). Since there are at most n possible values any $f(x)$ can take, and only k distinct x, the loop certainly executes not more than kn times. So, if we can establish that the computation in the loop can be carried out in constant time, the claim follows.

In an initialisation phase (not part of the algorithm proper) we can certainly compute tables of all the values z^a_b for z in both π and τ, $a \in \{\blacktriangleleft, \blacktriangleright, \blacktriangledown, \blacktriangle\}$ and b either absent or equal to one of L or U. For π this can be done in $O(k)$ time, and for τ in $O(n)$ time, so this initialisation can be absorbed into the claimed run-time. This ensures that the "Update" operations in the loop can be carried out in constant time. We can maintain $P(f)$ as a queue, and separately maintain an array of boolean values that indicate whether or not $x \in P(f)$. To start the loop, we dequeue some x. The update operation ensures that x is no longer a problem, so we can set its value in the array to `false`. Moreover, the update operation only changes the value of $f(x)$, and increases it. So it cannot "solve" any existing problem (other than that of x) and the only other way that it could change the problem set would be if $f(x)$ moved to the right of $f(x^{\blacktriangleright})$ or above $f(x^{\blacktriangle})$. Therefore, in the recompute phase we only need to check those two possibilities, and enqueue $x^{\blacktriangleright}$ and/or $x^{\blacktriangle}$ (setting their boolean values in the array to `true`) if necessary. By making reference to the array, we can ensure that we never have duplicate elements in the queue – so every iteration of the loop really does result in a proper update. □

3. FLUID ELEMENTS AND THE $O(kn)$ ALGORITHM FOR 321-AVOIDING TEXTS

In this section we aim to complete the proof of Theorem 1 and to do so we must face the issue of fluid elements in the pattern π. Since a fluid element participates in no 21, each other element of π is either below and left of it, or above and right of it. This is represented most easily using another notational convention. Suppose that σ and θ are two permutations of size m and n respectively. Then $\sigma \oplus \theta$ is the permutation whose points are:

$$\sigma \cup \{(i+m, \theta(i)+m) \,:\, i \in [n]\}.$$

Informally, to form $\sigma \oplus \theta$ we just place θ above and to the right of σ. Clearly $\oplus$ is associative, though of course not commutative.

For any 321-avoiding permutation π there is a unique representation:

$$\pi = \pi_1 \oplus \pi_2 \oplus \cdots \oplus \pi_t$$

where, for $1 \leqslant i \leqslant t$, π_i is either rigid or a singleton, and it is never the case that both π_i and π_{i+1} are rigid. The singleton elements of this representation correspond precisely to the fluid elements of π. For an example, consider again Figure 1 where the black squares correspond to the blocks π_i of this representation.

Given π of length k we can easily compute this representation in $O(k)$ time, simply by finding the fluid elements of π (which are those elements that are both left-to-right maxima and right-to-left minima). Henceforth, we assume that this representation is given.

In the algorithm to determine whether π embeds in τ we will construct, for each $1 \leqslant i \leqslant t$ at most two embeddings of $\pi_1 \oplus \cdots \oplus \pi_i$ into τ in such a way that, if any embedding of π into τ exists, then at least one of the two partial embeddings can be extended to a full embedding.

So we first consider the following question: given an embedding, e_i, of $\pi_1 \oplus \cdots \oplus \pi_i$ into τ that extends to an (unknown) embedding, e, of π into τ, how can we construct a pair of embeddings of $\pi_1 \oplus \cdots \oplus \pi_i \oplus \pi_{i+1}$ into τ, at least one of which extends to an embedding of π into τ?

We distinguish three cases for π_{i+1}. For this purpose, let T_i denote the set of elements that lie above and to the right of the image of e_i. Then, the image of e restricted to the elements corresponding to π_{i+1} is contained in T_i. Let us first consider the case where π_{i+1} is rigid. Then the image of e on the elements corresponding to π_{i+1} must be greater than or equal to (point by point), the image of π_{i+1} under its minimum embedding into T_i. Thus, if we choose the minimal embedding of π_{i+1} into T_i, the resulting embedding e_{i+1} extends to an embedding of π into τ. Though T_i is, strictly speaking, not a permutation all of its associated operators are the same as those of τ (except some are undefined, e.g., the leftmost element of T_i has no left neighbour in T_i but may well have one in τ). So, in this case we can use a slightly modified version of Algorithm 1 to compute a single extension of e_i with the required property.

A similarly easy case is where π_{i+1} is a singleton, i.e., a fluid element and T_i begins with its least element (which is a fluid element as well). Then nothing can be lost by mapping π_{i+1} to that element.

The only remaining case is where π_{i+1} is a singleton and the first element of T_i is not its minimum. Since every element of T_i lies above its first element, or above and to the right of its minimum,

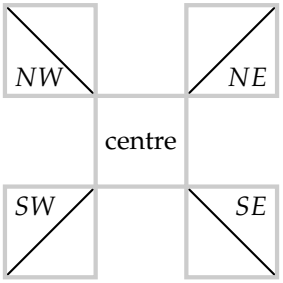

Figure 2: The decomposition of a skew-merged permutation into its centre and four corners.

we can extend e_i in two ways – one sending π_{i+1} to the leftmost element of T_i (which is an upper element) and one to its minimum (which is a lower element), and one of these must be extensible.

Now it seems that we might have a problem – given two partial embeddings of $\pi_1 \oplus \cdots \oplus \pi_i$ might they not extend to three or four candidate embeddings of $\pi_1 \oplus \cdots \oplus \pi_i \oplus \pi_{i+1}$? Indeed this is the case, but only if π_{i+1} is a singleton. If it has four possible images, two belong to U_τ and two to L_τ. Since all further elements of π lie above and to the right of this fluid element, we only need to retain the embeddings where its image is the lesser of the two in each of these sets. Likewise, if it has three possible images (one of which might be fluid), at least one of them can be ignored. Another way to say this is that because $\pi_1 \oplus \cdots \oplus \pi_i \oplus \pi_{i+1}$ ends with its maximum element, so do its images under the embedding. Among three or more elements of a 321-avoiding permutation there are at least one and at most two elements that do not participate as the 2 in a 12 pattern. We need retain only those embeddings whose maximum is not such a 2, as otherwise they could be replaced by an embedding with a smaller maximum in forming a full extension.

Since the sum of the length of the rigid permutations in the representation of π is at most the total length of π, the parts of the algorithm where we construct minimal rigid embeddings still require at most $O(kn)$ time in total. Each of the singleton extensions and filtering where necessary can actually be carried out in constant time, so the total cost of the algorithm is still $O(kn)$. If τ contains π the algorithm terminates successfully and returns one or possibly two embeddings.

What if no embedding exists? Then, following the plan above as if it did (beginning from an empty map, i.e., the case $i = 0$) we must at some point reach a failing case of Algorithm 1, or possibly encounter an empty T_i. In either case, we fail since we have demonstrated that no embedding can be possible.

This completes the proof of Theorem 1.

4. SKEW-MERGED PERMUTATIONS

The permutations avoiding 321 can be partioned into two monotone increasing sequences. Of course the permutations avoiding 123 can be similarly partitioned (into decreasing sequences) and the results of the previous section apply to them as well. However, the class of *skew-merged permutations*, those that can be partitioned into an increasing and a decreasing sequence, requires further analysis, though as we shall see the analogue of Theorem 1 is also true in this context.

Towards this goal, we first identify a set of *rigid* elements of a skew-merged permutation. In Figure 2 these are the elements lying in the corner regions. Specifically we say that an element of a skew-merged permutation is of type:

NE if it participates as a 3 in a 213;
NW if it participates as a 3 in a 312;
SW if it participates as a 1 in a 132;
SE if it participates as a 1 in a 231,

and we call any other element of a skew-merged permutation *central*. We first verify that the illustration of a skew-merged permutation shown in Figure 2 is correct. This is a result due to Atkinson [4], and so we only sketch part of the proof to give its flavour.

Proposition 7. *The elements of a skew-merged permutation decompose by type as shown in Figure 2. Moreover, the central elements form a monotone subsequence.*

Proof. Recall that another characterisation of skew-merged permutations is the following: they are those permutations that do not contain either 3412 or 2143.

Let a skew-merged permutation π be given, and suppose that $\pi = I \cup D$ is a partition of π into a monotone increasing and monotone decreasing sequence. Consider first elements of type NE (all other types can be handled by parallel arguments due to symmetry). Since any such participates as a 3 in a 213, it must belong to I (otherwise, the elements participating as the 2 and 1 would both belong to I which is of course impossible). So the elements of type NE form a monotone increasing sequence.

Suppose that C is of type NE, with BAC an occurrence of 213 and a is of type SW with acb an occurrence of 132. Then $a \in I$ for similar reasons to the preceding ones. If C preceded a (and hence was also smaller than it) then, $BAcb$ would be an occurrence of 2143. So, all elements of type SW lie below and to the left of those of type NE. Now suppose that z is of type NW, with zxy an occurrence of 312. If C were to precede z we would have various cases: first if C lay below y then $BAzy$ would be 2143, if C lay above y but below z then $Czxy$ would be 3412, if C lay above z and B above y, then $BCxy$ would be 3412, but if B lay below y then $BAzy$ would be 2143. As all these cases lead to contradictions, C must follow z.

All other cases can be dealt with similarly. Finally, to see that the central elements form a monotone sequence observe that they must certainly avoid all of 132, 213, 231, and 312 lest some of them be non-central. But, only monotone permutations (of either type) avoid these four permutations. □

This decomposition can be computed in linear time:

Lemma 8. *Given a skew-merged permutation of length n, there is an algorithm that computes its partition into types in $O(n)$-time.*

Proof. Let θ be an arbitrary skew-merged permutation. Notice that the part of θ to the left of the leftmost element of type NE or SE avoids 231 and 213. Such permutations have a characteristic $>$ shape since any element must not be intermediate in value between two to its right. We are interested in finding the maximum prefix of θ which has this characteristic shape, or what amounts to the same thing, the leftmost element of θ such that the prefix ending at that element involves 231 or 213. This is easily accomplished in linear time: simply scan θ from left to right continuing so long as the prefix is monotone. Once an element is encountered that violates monotonicity, it, and

all subsequent elements must have values that lie between the values of the two immediately preceding elements. The first time this condition fails (if it ever does), we have identified the leftmost element of type NE or SE. That is, we have found the boundary line between the centre region and the Eastern region of θ.

In a similar manner we can find all of the boundary lines: by scanning θ from right to left we find the boundary between West and centre, by scanning from bottom to top we find the boundary between South and centre and by scanning from top to bottom we find the boundary between North and centre. We can thus compute the partition of θ into types by scanning θ four times. □

The non-central elements of a skew-merged permutation correspond to the rigid elements of a 321-avoiding permutation. Since they are defined by the occurrence of certain patterns and since embeddings preserve such patterns it is immediately clear that if $e : \pi \to \tau$ is an embedding of one skew-merged permutation into another, then e must preserve the type of all non-central elements.

For two non-central elements of the same type we define their minimum to be the one further from the centre. For two embeddings, e_1 and e_2 of the non-central elements of π into the non-central elements of τ (that preserve type) we can define their meet by

$$e_1 \wedge e_2(x) = \min\{e_1(x), e_2(x)\}$$

for all non-central $x \in \pi$. Then, just as in the 321-avoiding case it is easy to see that $e_1 \wedge e_2$ is also an embedding of the non-central elements of π into those of τ. Moreover, if either e_1 or e_2 was the restriction of an actual embedding, e, of π into τ to the non-central elements then we can extend the mapping $e_1 \wedge e_2$ to central elements using e there, and thereby obtain an embedding. So, among all embeddings of π into τ there is one whose effect on the non-central elements is the minimum of all the embeddings of the non-central elements of π into those of τ.

This minimum map on the non-central elements can be found by the obvious modification of Algorithm 1. The only slight complications are in the form of the update rule (since there are four types of elements and not just two), and the recomputation of the problem set (for similar reasons). However, given partitions of both π and τ by type (and tables of predecessors and successors by position and value both globally and by type), both update and recomputation steps can still be carried out in constant time.

To complete a proof of Theorem 2 we must show that these tables can be computed sufficiently quickly, and that, having found a minimum embedding of the non-central elements of π to those of τ, the existence (or otherwise) of a full embedding can also be determined sufficiently quickly.

To compute the various tables is easy. For instance for the successor by position for elements of type SE we need only compute (both) the successor of type NE and the successor of type SE. But this is easily accomplished in a single pass through these elements (which form an initial segment by position of the whole permutation). Each additional element is the successor of the preceding element of its type (within its type), and of all the elements of the other type which do not yet have successors assigned (elements waiting for a successor can be added to a queue as they are encountered and removed en masse when the successor is found.)

Having found the minimum embedding of the non-central elements of π into τ (if no such embedding exists, we need continue no further), it remains to determine whether or not the central part of π can be embedded into the remainder of τ (which may include both central, and non-central ele-

ments). The central part of π is a monotone pattern of a certain length at most k, and the remaining part of τ is a skew-merged permutation of length at most n (whose endpoints we know).

In general, finding a longest increasing (or decreasing) subsequence of length k in a permutation of length n can be done in time $O(n \log \log k)$ [7]. Thus, checking whether the central part can be embedded into the remaining part of τ can be done within the $O(nk)$ runtime bound of our algorithm. In the special case of skew-merged permutations, finding a longest increasing (resp., decreasing) subsequence can even be done in $O(n)$ time. To be more precise, $O(n)$ time is only required for obtaining the partition into five types as represented in Figure 2 (which is already available in our case); the remaining steps require only constant time.

Indeed, for longest increasing subsequences the following observations can be made (the case of decreasing subsequences can be treated analogously): The elements of type SW and NE will always contribute to a longest increasing subsequence. Moreover, such a subsequence also contains as many elements as possible from the centre, i.e., if the center is increasing then all elements contribute to a longest increasing subsequence and if the center is decreasing we can arbitrarily pick one centre element. Note that it is never advantageous to include elements of type NW or SE. This can be seen as follows: At most one NW or SE element can be part of an increasing subsequence. Thus, if the centre is non-empty, it is certainly not advantageous to include a NW or SE element. Let us assume that the centre is empty. An element of type NW occurs as a 3 in a 312 pattern. Among the elements playing the role of the 1 and the 2, at least one element (and possibly both of them) is of type SW or NE. Thus, including an element of type NW would force us to exclude one or two elements of type SW or NE. In other words, we cannot increase the length of an increasing subsequence by adding an element of type NW. A similar argument holds for elements of type SE. We conclude that for the length of the longest increasing subsequence only the number of elements of type SW and NE as well as the number of increasing central elements are relevant. Given a decomposition into types, computing these numbers can clearly be done in constant time.

5. CONCLUDING REMARKS

We conclude by mentioning some open problems related to this work. We have seen in Theorem 1 that Av(321)-PPM can be solved in $O(kn)$ time. Guillemot and Vialette showed that the more general Av(321)-PATTERN PPM problem can be solved in $\mathcal{O}(kn^{4\sqrt{k}+12})$ time. It is an open problem whether Av(321)-PATTERN PPM can be solved in polynomial time. Note that if the pattern avoids 132, 231, 213 or 312 then it is automatically separable and thus the $\mathcal{C}$-PPM problem for all four classes Av(132), Av(231), Av(213) or Av(312) can be solved in polynomial time. Consequently the Av(321)-PATTERN PPM—which is equivalent to the Av(123)-PATTERN PPM—is the only open case for Av(β)-PATTERN PPM where β has length 3.

In case Av(321)-PATTERN PPM turns out to be NP-complete, Av(β)-PATTERN PPM will also be NP-complete if β is any permutation of length four other than 2143, 3142, 2413, or 3412. Interestingly, this list contains exactly those patterns that define the classes of skew-merged and of separable permutations. Moreover, NP-completeness of Av(321)-PATTERN PPM would imply that Av(β)-PATTERN PPM is NP-complete for β of length five or more, since by Erdős–Szekeres Theorem [8] every permutation of length five contains 123 or 321.

Looking at the big picture, Theorems 1 and 2 show that $\mathcal{C}$-PPM can be solved in polynomial time

for $\mathrm{Av}(321)$ and $\mathrm{Av}(2143, 3412)$, respectively. It might be that $\mathcal{C}$-PPM is always polynomial-time solvable for a fixed, proper class $\mathcal{C}$. It would be of considerable interest to either establish this statement or to prove a dichotomy theorem that distinguishes permutation classes for which $\mathcal{C}$-PPM is polynomial-time solvable and those that yield hard $\mathcal{C}$-PPM instances. The same question can be asked for $\mathcal{C}$-PATTERN PPM, although it seems rather unlikely that this problem is polynomial time solvable for every fixed, proper class $\mathcal{C}$.

REFERENCES

[1] AHAL, S., AND RABINOVICH, Y. On complexity of the subpattern problem. *SIAM J. Discrete Math. 22*, 2 (2008), 629–649.

[2] ALBERT, M., ALDRED, R., ATKINSON, M., AND HOLTON, D. Algorithms for pattern involvement in permutations. In *Algorithms and Computation*, P. Eades and T. Takaoka, Eds., vol. 2223 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, 2001, pp. 355–367.

[3] ALBERT, M. H., ATKINSON, M. D., BRIGNALL, R., RUŠKUC, N., SMITH, R., AND WEST, J. Growth rates for subclasses of $\mathrm{Av}(321)$. *Electron. J. Combin. 17* (2010), Paper 141, 16 pp.

[4] ATKINSON, M. D. Permutations which are the union of an increasing and a decreasing subsequence. *Electron. J. Combin. 5* (1998), Paper 6, 13 pp.

[5] BOSE, P., BUSS, J. F., AND LUBIW, A. Pattern matching for permutations. *Inform. Process. Lett. 65*, 5 (1998), 277–283.

[6] BRUNER, M.-L., AND LACKNER, M. A fast algorithm for permutation pattern matching based on alternating runs. *Algorithmica* (2015), 1–34.

[7] CROCHEMORE, M., AND PORAT, E. Fast computation of a longest increasing subsequence and application. *Information and Computation 208*, 9 (2010), 1054 – 1059.

[8] ERDŐS, P., AND SZEKERES, G. A combinatorial problem in geometry. *Compos. Math. 2* (1935), 463–470.

[9] GUILLEMOT, S., AND MARX, D. *Finding small patterns in permutations in linear time*. 2014, ch. 7, pp. 82–101.

[10] GUILLEMOT, S., AND VIALETTE, S. Pattern matching for 321-avoiding permutations. In *Algorithms and Computation (20th International Symposium, ISAAC 2009, Honolulu, Hawaii, USA, December 16-18, 2009)* (Berlin, 2009), Y. Dong, D.-Z. Du, and O. Ibarra, Eds., vol. 5878 of *Lecture Notes in Comput. Sci.*, Springer, pp. 1064–1073.

[11] IBARRA, L. Finding pattern matchings for permutations. *Information Processing Letters 61*, 6 (1997), 293–295.

[12] SAXENA, S., AND YUGANDHAR, V. Parallel algorithms for separable permutations. *Discrete Applied Mathematics 146*, 3 (2005), 343–364.

[13] STANKOVA, Z. E. Forbidden subsequences. *Discrete Math. 132*, 1-3 (1994), 291–316.